
NORMATIVITY AND INTENTIONALITY

DRAFT

Anandi Hattiangadi

The hard problem of intentionality is to explain what makes it the case that an arbitrary sentence or thought has the semantic properties that it does rather than some other semantic properties or none at all. Some hold that intentionality is normative, and that this has a crucial bearing on the hard problem of intentionality. This paper investigates whether this is so.

It is possible to distinguish four versions of the thesis that intentionality is normative: a) grasp of a concept or meaning involves following a rule or making a normative judgment of some kind; b) the concepts of meaning and content are normative; c) meaning and content are sources of normativity; d) the semantic facts are in some sense reducible to normative (and natural) facts. I discuss all four versions of the thesis, and argue that the normativity of intentionality has little bearing on the hard problem of intentionality.

1 INTRODUCTION

Philosophy, linguistics, and psychology have furnished us with a vastly improved understanding of intentionality, the capacity to represent something or other in some

way. Yet the so-called ‘hard problem’ of intentionality remains unresolved.¹ This is the problem of explaining how and why the semantic properties of intentional states—their meaning or content²—arise in the world as described by the natural sciences, such as most notably, physics, chemistry and biology. Since intentionality is unlike anything else we encounter in the natural world, it gives rise to some hard questions: can the semantic truths be reductively explained in terms of the natural truths, or more broadly, in terms of the non-semantic truths? Do the semantic facts supervene³ on the non-semantic facts? What makes it the case that an arbitrary representation has the semantic properties that it does rather than some other semantic properties or none at all?

This essay focuses on approaches to the hard problem of intentionality that in some way or other appeal to the idea that intentionality is *normative*. This idea took shape in Saul Kripke’s influential discussion of Wittgenstein’s rule following considerations. Recall that Kripke asks what makes it the case that he means addition by ‘plus’ rather than quaddition (where x quus $y = x$ plus y if $x, y \leq 57$, and $= 5$,

¹ Chalmers 1996 distinguishes the hard problem of consciousness, which is to reductively explain consciousness in non-phenomenal terms, from the easy problems, having to do with such things as finding correlations between conscious experience and states of the brain, or with understanding the relationships between different kinds of conscious experience. For Chalmers, the distinction between the hard problems and the easy problems has to do with the methods by which they can be resolved: the easy problems can be resolved by ordinary empirical methods, whereas the hard problems require a priori reasoning. This distinction between methods appropriate to the resolution of hard and easy problems is not directly relevant to the present discussion.

² Typically, meaning is said to be a property of sentences, expressions or utterances, while content is said to be a property of mental states, such as beliefs, desires, intentions, concepts or mental representations. It may be that there is a kind of intentionality that does not involve meaning or content. We can set aside that kind of intentionality here.

³ Roughly, to say that the A-facts supervene on the B-facts is to say that any two things that are qualitative duplicates in all their B-properties are duplicates in all their A-properties. There are many varieties of supervenience. For a survey of them, see CITE.

otherwise), and reaches the skeptical conclusion that there is no fact of the matter what any word means. He rejects the view that meaning is determined by speaker dispositions on the grounds that ‘the relation of meaning and intention to future action is *normative*, not *descriptive* (Kripke 1982: 37).’ The problem, Kripke says, is ‘that my present mental state does not appear to determine what I *ought* to do in the future (Kripke 1982: 56).’ Similarly, in response to the view that meaning is determined by an experience with a distinctive qualitative feel, he asks, ‘How on earth would this [experience] help me figure out whether I ought to answer ‘125’ or ‘5’ when asked about ‘68+57’? If the [experience] indicates that I ought to say ‘125’, would there be anything about it to refute a sceptic’s contention that, on the contrary, it indicates that I should say ‘5’ (Kripke 1982: 42)?’ Though Kripke directly discusses the normativity of meaning, it is clear that his claims must generalize, and he is committed to the normativity of content as well (Boghossin 1989, Hattiangadi 2007).

Inspired in part by Kripke, several philosophers have argued that meaning is normative (Boghossian 1989, 2005 Brandom 1994, 2000, Gibbard 1994, 2013, Ginsborg 2011, 2012, Whiting 2007), content is normative (Boghossian 2003, Gibbard 1996, 2003, Wedgwood, 2007), or more generally that intentionality is normative (Wedgwood 2007). One underlying motivation for this is the thought that there is a useful analogy to be drawn between the hard problem of intentionality and the hard problem of normativity—the problem of explaining what makes it the case that something is right, wrong, good, bad, obligatory, forbidden or permitted in non-normative terms. This certainly seems to be Kripke’s motivation. Though one might reject Kripke’s anti-realism about normativity, and resist his skeptical conclusion about meaning, it is worth considering whether the normativity of meaning or content has any bearing on the hard problem of intentionality. If it does, then well-developed

approaches to the hard problem of normativity might be fruitfully carried over to the hard problem of intentionality.

What exactly does it mean to say that intentionality is normative? First of all, it will help to distinguish between normativity as a feature of a certain class of representations and normativity as a feature of the world. It is relatively uncontroversial that there are normative judgments, statements, concepts and expressions, such as the statement ‘FGM ought to be prohibited’, the judgment that all and only pleasure is good, the expression ‘good’ as used in ‘all and only pleasure is good’ and the concept ‘good’ is used to express. More controversial is the thought that in addition to normative representations, there are normative facts and properties to which normative representations correspond, such as the fact that FGM ought to be prohibited or the property of goodness. Those who claim that intentionality is normative can be viewed as using ‘normative’ in either of these two ways. Indeed, we can distinguish four hypotheses—two of each kind.

1. *Meaning or content essentially involves normative judgment.* This is the view that what it is to grasp a concept or a meaning involves following a constitutive rule, or more generally, that grasping a meaning or concept involves making a normative judgment of some kind. For instance, a proponent of this view might say that if Shah means addition by ‘plus’, then she follows a rule for addition, or is disposed to judge that ‘ $57+68=125$ ’ is correct.
2. *Semantic concepts are normative.* According to this view, the *concept* ‘meaning’ is normative, as are straightforward ascriptions of meaning and

content. For instance, a proponent of this view might say that the statement ‘Shah means addition by “plus”’ is a normative statement and the judgment that Shah believes that $57+68=125$ is a normative judgment.⁴

3. *Meaning and content are sources of normativity.* This is the view that semantic facts or properties determine or give rise to normative facts or properties. For instance a proponent of this view might say that necessarily, if Shah means addition by ‘plus’, then she ought to accept that $57+68=125$, or at least that she has a normative reason to do so.^{5,6}
4. *Semantic facts are reducible to the normative (and natural) facts.* According to this view, normative facts, perhaps together with the physical facts, determine the semantic facts. For instance, a proponent of this hypothesis might say that necessarily, if Shah ought to accept a sentence s iff p is true, then s means that p in the language Shah speaks.

These four versions of the hypothesis that intentionality is normative are not mutually exclusive. Indeed, they are often held in tandem. But since they are independent, it will make sense to consider them separately. Is intentionality normative in any of these four senses? And what does the normativity of intentionality in any of these four senses have to do with the hard problem of intentionality? As we shall see, normativity ultimately has very little bearing on the hard problem of intentionality.

⁴ Not all such statements and judgments, however. For instance, sometimes statements involving normative terms are used in the ‘inverted commas sense’, that is, as an implicit quotation.

⁵ This is what Kathrin Glüer and Åsa Wikforss call ME/CE normativism, for meaning-engendered and content-engendered norms.

⁶ I assume that to accept a sentence is either to assert it, assent to it, or to be disposed to do so.

2 MEANING, CONTENT AND NORMATIVE JUDGMENT

Kripke introduces the thought that meaning is normative in the context of a discussion of Wittgenstein's rule following considerations. At the heart of these reflections on rule following is the idea that language is a conventional activity, which invites an analogy between language and other conventional activities such as games (Wittgenstein 1953): just as the activity of playing bridge is in part constituted by the rules of bridge, so too is the activity of speaking a language. Kripke suggests that meaning essentially involves rule following, as when he says that '[o]rdinarily, I suppose that, in computing '68+57' as I do, I do not simply make an unjustified leap in the dark. I follow directions I previously gave myself that uniquely determine that in this new instance I should say '125' (1982:CITE),' and that such directions 'must somehow be "contained" in any candidate for the fact as to what I meant (1982:CITE).'

The hypothesis that speaking a language involves rule following is a special instance of the more general thesis that meaning something by an expression, or grasping a concept, essentially involves making a normative judgment of some kind. As is widely agreed, following a rule involves more than merely acting in accordance with it; it involves in some sense being guided by or accepting the rule. And this involves a normative judgment. Even if the idea that meaning involves rule-following is waning in popularity, many philosophers are drawn to the more general thesis that meaning something or grasping a concept essentially involves a normative judgment of some kind. For instance, normative inferentialists, such as Robert Brandom, hold that for a speaker to grasp the meaning of a sentence she must participate in a practice by accepting certain inferential rules, such as a rule that permits inference from 'roses

are red' to 'roses are coloured', and 'entrance' rules, which specify the conditions under which one is permitted to assert a sentence, such as when one has suitable, undefeated perceptual evidence for its truth (Brandom 1994). Similarly, conceptual role semanticists, such as Christopher Peacocke (1992) and Ralph Wedgwood (2009), hold that grasping a concept such as 'if' involves normative commitments and normative judgments. For Wedgwood, grasp of the concept 'if' involves a rational disposition to accept inferences by modus ponens: 'if A then B; A; therefore B' (Wedgwood 2009). A rational disposition is a disposition that is sensitive to the rationality of a transition in thought. To grasp the concept 'if', it is not enough to be disposed accept inferences by modus ponens; it must be the rationality of those inferences that trigger the disposition. Thus, grasp of a concept involves at least an implicit judgment that certain patterns of inference are rational. Assuming that judgments of what is rational are normative, it follows that grasp of a concept involves a normative judgment of some kind (See also Boghossian 2003, Ichikawa and Jarvis 2013).

Similarly, Hannah Ginsborg defends a view of this general kind, though she claims that meaning essentially involves implicit rule following rather than explicit rule following. What it is to implicitly follow a rule for Ginsborg is to be disposed to make the appropriate pattern of normative judgments. For instance, on her view, it is necessary for meaning addition by 'plus' that a subject be disposed to judge that '57+68=125' is correct, that '125' is the answer one ought to give to '68+57=?'. She describes this as a 'primitive' attitude in the sense that it need not be the result of conscious rule following, or even the product of a post hoc realization of which rule one has been following (Ginsborg 2011).

Does intentionality involve normative judgment, as these philosophers suggest? There are some excellent reasons to think not. First of all, rules that might be thought to be constitutive of meaning are not the sorts of rules that can be followed (Glüer & Pagin 1999, Glüer 1999). Constitutive rules of meaning specify what a sentence means, and this does not naturally coincide with telling speakers what to do. Consider, for instance, a rule that says: ‘plus’ means addition in English. This is a *descriptive* rule, which captures a generalization about usage. Since it does not tell anyone what to do, it is not at all obvious what is involved in following it. Indeed, even if you accept this rule, your acceptance of it can only play the role of a belief in practical reasoning (Cf. Glüer and Pagin 1999, Glüer & Wikforss 2010).

Some have sought to sidestep this problem by formulating constitutive rules in terms of correctness conditions, and then arguing that correctness is a normative notion (Boghossian 1989, 2003, Gibbard 2003). The rule for meaning addition by ‘+’ might then be stated as follows:

- (1) ‘ $\varphi + \psi = \omega$ ’ is correct iff the number denoted by ω is the sum of the numbers denoted by φ and ψ .

The crucial claim is that ‘correct’ is a normative term, having to do with what one ought to do or may do (Boghossian 2003, Gibbard 2003). So, to accept (1) is to make a normative judgment. But ‘correct’ is not a normative term across all contexts (Bykvist & Hattiangadi 2013), and it is far from obvious what the normative implications of (1) are supposed to be. If I accept (1), do I judge that I ought to assert every addition fact,

or merely that I may do so? Do I judge that I ought to blurt out addition facts at any time, or only when I am asked? A natural reading of (1) is that in order to mean addition by 'plus' you need to take yourself to be under an obligation to respond with the sum of any two numbers when asked. But this view is implausible. Suppose that you discover out that if you answer with the sum of two numbers in some situation, thousands will die. Though you will undoubtedly judge in this situation that you ought not to respond with the sum, we can hardly conclude that you do not mean addition by 'plus'. Similar considerations tell against the interpretation of (1) according to which in order to mean addition by 'plus' you must judge that you *may* respond with the sum of any two numbers when asked (E.g. Whiting 2009). For in the circumstance just described, you sensibly judge that you ought not to respond with the sum, which is incompatible with the judgment that you may do so. Nearly as implausible is the thought that meaning addition by 'plus' requires that you judge yourself to have a *pro tanto* reason to assert ' $57+68=125$ ', that is, a reason that can be outweighed by other considerations. However, an unrepentant liar might not take there to be any even *pro tanto* reason to assert an addition truth in some situation. Once again, we cannot conclude that the unrepentant liar does not mean addition by 'plus' (Cf. Glüer 1999, 2001, Hattiangadi 2006, 2007, Wikforss 2001).

Perhaps the conclusion to draw is that correctness is a *sui generis* normative notion, irreducible to obligation and permission (See McHugh 2013). One might then say that the liar must at least judge that responding with the sum is correct in view of what 'plus' means, even if she does not take herself to have any reason to do so. But it is difficult to see in what sense correctness is a normative notion if it is possible to coherently judge that some act is correct while simultaneously judging that one has no normative reason to perform that act. Even if correctness is irreducible to any other

normative notions, it must bear some relation to them. McHugh suggests that if one judges an act to be incorrect, then one judges that one has a *pro tanto* reason not to perform that act. But as the case of the unrepentant liar proves, it is possible for someone to think that they have no *pro tanto* reason not to respond with the sum in some situation without thereby failing to mean addition by ‘plus’.

Ginsborg puts a kind of particularist twist on the idea that meaning involves rule following. According to her, to mean something by an expression, a subject need not explicitly endorse any rule, but must be disposed to make the appropriate pattern of particular judgments of correctness. Yet the particularist twist does not really help to address the foregoing objections. Since judgments of correctness are not normative in every context, we need to know what the normative implications of these primitive judgments of correctness are. Perhaps Ginsborg will say that to mean addition by ‘plus’ you must be disposed to judge that you ought to answer with the sum of any two numbers when asked. But consider once again the situation in which you judge that you ought not to assert that $57+68=125$ in order to avoid the loss of life. Though you are not disposed to make the primitive judgment that is constitutive of meaning addition by ‘plus’, it does not follow that you do not mean addition by ‘plus’. Perhaps Ginsborg will instead say that to judge that some answer is correct is to judge that you have a normative reason to give that answer. But once again, an unrepentant liar who sees herself as having no reason to say that $57+68=125$ in some contexts, need not fail to mean addition by ‘plus’.

Perhaps the suggestion that grasping a concept essentially involves a normative judgment fares better. Consider the conceptual role semanticist’s suggestion that in order to grasp the concept ‘if’, a subject must have a rational disposition to accept arguments by *modus ponens*. We can put pressure on this

suggestion by appeal to Timothy Williamson's example of Vann McGee, a distinguished logician who has presented counterexamples to modus ponens, such as the following (Williamson 2007:92ff). Suppose that the opinion polls place the Republican Ronald Reagan decisively ahead of the Democrat Jimmy Carter, who in turn is placed decisively ahead of the other Republican, John Anderson. It is reasonable to believe that if a Republican wins the election, then if it is not Reagan who wins it will be Anderson. It is also reasonable to believe that a Republican will win the race. Yet McGee argues that it would not be reasonable to conclude that if it is not Reagan who wins, it will be Anderson.

In presenting this counterexample, McGee shows that he lacks the rational disposition to accept arguments by modus ponens in all cases, for in some cases, he judges it irrational to accept an argument by modus ponens. Yet, as a distinguished logician, it is unquestionable that McGee grasps the concept of the conditional; if anything, he is an expert on conditionals (Williamson 2007). So, it appears to be possible to grasp the concept of the conditional while rejecting as irrational some arguments by modus ponens.

Perhaps it will be said that though McGee does not accept all arguments by modus ponens, he accepts enough of them to come close enough to qualify as competent in his grasp of the concept of the conditional. The trouble with this reply is that though McGee may come close enough to qualify as competent with the concept of the conditional, he better qualifies as competent with a distinct concept—call it the concept of the McGee conditional. Let's say that to grasp the McGee conditional you have to be rationally disposed to accept inferences by McModus ponens: 'if A then B; A; therefore B, in all contexts except the McGee cases', where the McGee cases are those cases that McGee puts forward as counterexamples to modus ponens. McGee

better qualifies as grasping the concept of the McGee conditional than as grasping the conditional, for he accepts all inferences by McModus ponens, and this is constitutive of the concept of the McGee conditional. If what it is to grasp a concept is to have the rational dispositions constitutive of that concept, we should say that McGee grasps the concept of the McGee conditional not the conditional.

Could it be that McGee cannot have a *rational* disposition to accept arguments by McModus ponens, because that rule is not a requirement of rationality? If McModus ponens is not a rational rule, it is not possible to have a disposition to accept arguments by McModus ponens because those arguments are rational. However, this reply has the implausible consequence that it is impossible for anyone to grasp the concept of the McGee conditional. Though McGee himself surely grasps the concept of the conditional, it is possible that someone with McGee's dispositions grasps the concept of the McGee conditional. But if it is a necessary truth that one rationally ought to accept arguments by modus ponens, then it is not possible for anyone to have a rational disposition to accept arguments by McModus ponens, and it is not possible for anyone to grasp the concept of the McGee conditional. But this consequence is unacceptably strong. Though admittedly bizarre, the concept of the McGee conditional is a *bona fide* concept. It is easy enough to grasp the concept whether or not one accepts that McGee's cases really constitute counterexamples to modus ponens.

It is time to move on to the question whether the hypothesis that meaning and content involve normative judgments of some kind has any bearing on the hard problem of intentionality. Kripke certainly suggests that it does. He takes the upshot to be that there are no semantic facts, 'no such thing as meaning anything by any word'. Similarly, both Brandom and Ginsborg present themselves as addressing

Kripke's problem and proposing an alternative solution. Perhaps these authors have in mind a different version of the thesis that intentionality is normative. Because even if it is true that meaning and content involve normative judgments of some kind, this ultimately has no bearing on the hard problem of intentionality.

It is important to distinguish the hard problem of intentionality from what might be called 'first order' semantics.⁷ The principle aim of first order semantics is to accurately model the intentional and semantic phenomena. Among other things, it concerns the relations between different kinds of intentional and semantic phenomena, including relations between the meanings of expressions and the meanings of sentences, and relations of priority or determination between different kinds of intentional state. For instance, it is a question of first order semantics whether sense determines reference, or whether reference is determined by causal-historical relations to an initial dubbing event. It is easy to confuse this first order determination relation with meta-semantic determination.⁸ What makes the causal-historical theory of reference a first-order theory is that the postulated determination relation holds between two kinds of *intentional* phenomena: contemporary uses of a referring expression and a prior intentional act of naming or stipulation, the 'dubbing event'. In the dubbing event, the intentions of the dubber play a crucial role in fixing the referent of the expression.⁹

⁷ This roughly corresponds to what Stalnaker (CITE) calls descriptive semantics, but the use of the term 'descriptive' in this context might be misleading.

⁸ Cf. Stalnaker (CITE). My reason for this is that the causal-historical theory relates current use of a referring expression to a prior intentional act—the dubbing event. In contrast, a causal theory of reference in the

⁹ In contrast, Fodor's (CITE) and Dretske's (CITE) theories of mental representation fall squarely in the meta-semantics. Though these theories also appeal to causal relations, they aim to be fully reductive.

Now, it is clearly a question of first order semantics whether meaning something by an expression or grasping a concept essentially involves making a normative judgment of some kind. This is because the question whether meaning or content involves normative judgment has to do with relations between intentional mental states, between different kinds of intentional or semantic phenomena—between the intentional mental state of meaning something by an expression or grasping a concept and the intentional mental state of following, endorsing or accepting a rule, or making a normative judgment of correctness. Construed as an answer to a first-order question concerning the interrelations between semantic phenomena, the hypothesis that meaning or content involves normative judgment might well be illuminating. But construed as an answer to the hard problem of intentionality, it is blatantly circular. The hard problem of intentionality is to reductively explain intentionality in non-semantic, non-intentional terms, and any unreduced appeal to normative judgments in the reductive explanation of intentionality would render the explanation circular. Of course, a dualist with respect to intentionality will argue that the hard problem of intentionality cannot be solved. But it is hard to see how the hypothesis that meaning and intentionality involve normative judgments would speak in favour of dualism. Boghossian (1989) suggests that the normativity of meaning and content figure in an argument for dualism. However, this conclusion can only be drawn if some other version of the normativity thesis is true, such as that semantic requirements are sources of normativity, or if normative truths figure essentially in the reductive explanation of intentionality.

Some have suggested that the hypothesis that meaning or content involve normative judgment constitutes a *partial* solution to the hard problem of intentionality (Ginsborg 2011). The idea might be that the explanation of intentionality proceeds in

stages: first, we reductively explain meaning and content in terms of normative judgments, and this paves the way for a reductive explanation of the contents of normative judgments. Ginsborg suggests that the contents of these normative judgments can be reductively explained in terms of dispositions, and that Kripke's objections to the dispositional theory can be met (though she doesn't herself say how). But if Kripke's objections to the dispositional theory can be met, a dispositional explanation of the contents of non-normative judgments is available and it is possible to avoid the step of explaining meaning in terms of normative judgments. Thus, to be told that meaning or content involve rule following is not immediately relevant to the hard problem of intentionality.

3 THE NORMATIVITY OF SEMANTIC CONCEPTS

In his recent book, Allan Gibbard argues not that meaning is normative, but that the concept MEANING is normative,¹⁰ much like paradigmatic normative concepts such as 'good' and 'ought'. He claims further that meta-representational ascriptions of meaning, such as 'Shah means addition by "plus"' are normative, not descriptive, as are self-ascriptions of meaning, and ascriptions of propositional attitudes. Gibbard combines this view with his own brand of expressivism about normativity according to which assertions of normative statements do not purport to describe how things are, but express something akin to plans. Gibbard proposes to shed light on the hard problem of intentionality by exploiting the analogy with meta-ethics. In meta-ethics, the expressivist says that because moral sentences are non-descriptive, there is no need to postulate any moral facts out there in the world to correspond to them, and *a*

¹⁰ For the duration of this section, I will adopt Gibbard's convention of using small capital letters to denote concepts and thoughts.

fortiori, no moral facts crying out for reductive explanation. By the same token, if semantic ascriptions express plans, there is no need to postulate any semantic facts out there in the world to correspond to them, and *a fortiori*, no semantic facts crying out for reductive explanation.

Gibbard introduces his view with the slogan ‘means entails ought’. ‘Ascriptions of meaning,’ Gibbard says, ‘imply straight ought ascriptions’ (Gibbard 2012: 11). He adds that these entailments ‘could only be analytic or conceptual’ (Gibbard 2012:23). That is, ascriptions of meaning, such as ‘Pierre means DOG by “chien”’, analytically entail ascriptions of oughts, such as ‘Pierre ought to...’; these entailments hold in virtue of the meaning of ‘meaning’ or equivalently, in virtue of the concept MEANING. Anyone who grasps the concept MEANING must accept the normative entailments of meaning ascriptions.

What are the normative entailments of meaning ascriptions? Here are two examples drawn from Gibbard’s discussion, where the first statement of each pair is said to analytically entail the second:¹¹

- 1a.** The sentence ‘schnee ist weiss’ in Ursula’s language means SNOW IS WHITE.

¹¹ The first example occurs on pg. 40 of Gibbard 2012, with minor stylistic variations, while the second is reconstructed from Gibbard’s discussion of analyticity and synonymy in chapter 6. Some further examples occurring in the text are somewhat puzzling. For instance, Gibbard says: ‘I ought not to believe, all at once, that snow is white and that nothing is white—and that ties in with the meaning of our term “nothing”’ (Gibbard 2012:13). Even if it is true that I ought not to believe both that snow is white and that nothing is white, this seems not to have anything to do with the meaning of the term ‘nothing’. For it is arguably true that Kaveri, a monolingual speaker of Konkani, ought to not to believe both that snow is white and that nothing is white, yet this has little to do with the meaning of our term ‘nothing’.

1b. Ursula ought to accept ‘schnee ist weiss’ iff snow is white.

2a. Pierre’s sentence ‘les chiens aboient’ means DOGS BARK.

2b. If Pierre has sufficient undefeated evidence that dogs bark in an epistemic circumstance *E*, then Pierre ought to accept the sentence ‘les chiens aboient’ in *E*.

If these entailments are analytic as Gibbard suggests, then anyone who grasps the concept MEANING must accept that **1a** entails **1b**, and that **2a** entails **2b**, on pain of irrationality or conceptual confusion. However, it is possible for someone to sensibly deny these entailments without thereby displaying either irrationality or confusion about the concept MEANING.

First, consider an evidentialist, such as William Clifford, who holds that one ought never to believe anything on insufficient evidence (Clifford 1877). Suppose that we present Clifford with a hypothetical scenario in which snow is in fact white; Ursula is a speaker of German; but Ursula has never seen snow, and lacks any testimonial evidence as to its colour. Since Ursula lacks sufficient evidence that snow is white, Clifford would judge that Ursula ought not to believe SNOW IS WHITE. Furthermore, according to Gibbard, one way to believe SNOW IS WHITE is to accept a sentence in a language one understands that means that snow is white (Cf. Gibbard 2012:27, 2012:44). So, Clifford’s judgment that Ursula ought not to believe SNOW IS WHITE is tantamount to the judgment that Ursula ought not to accept the sentence ‘schnee ist weiss’. Thus, Clifford would accept **1a** but not **1b**. Yet, in so doing, he

would display neither irrationality nor confusion about the concept MEANING. Even if we think Clifford is mistaken, his mistake does not result from a failure to grasp the concept MEANING, because his reasons for rejecting the entailment from **1a** to **1b** have to do with his acceptance of evidentialism, a substantive view in epistemology. Thus, **1a** does not analytically entail **1b** (Cf. Whiting 2015).

It can be shown in a similar fashion that the entailment from **2a** to **2b** is not analytic. Consider Blaise Pascal (1670), who holds that there can be pragmatic reasons for belief. Suppose that we present Pascal with a hypothetical situation in which Pierre is a monolingual speaker of French who has sufficient undefeated evidence that dogs bark, but is promised eternal bliss if he does not accept the sentence ‘les chiens aboient’ and eternal torture if he does. With regard to this situation, Pascal would accept **2a**, but deny that Pierre ought to believe DOGS BARK. Since Gibbard assumes that believing DOGS BARK is tantamount to accepting a sentence in one’s language that means DOGS BARK, it follows that Pascal would reject **2b**. Yet in so doing, he need neither display irrationality nor confusion about the concept MEANING. Even if Pascal’s view that there can be pragmatic reasons for belief is mistaken, his mistake is not conceptual. His grounds for rejecting the inference from **2a** to **2b** have to do with his pragmatism, not with his grasp of the concept MEANING.¹² Thus, **2a** does not analytically entail **2b**.

One way to respond to these objections might be to appeal to the notion that there are distinct spheres of normativity: such as *inter alia*, epistemic, prudential and

¹² One might try to argue that he fails to grasp the concept of belief, because beliefs cannot be formed at will, but that would not help Gibbard in his efforts to argue that the concept of meaning is normative. In any case, Gibbard accepts that beliefs can be formed at will (Gibbard 2012).

semantic normativity. One might then say that what **1a** and **2a** entail are **1c** and **2c**, respectively:

1c. Ursula *semantically* ought to accept ‘schnee ist weiss’ iff snow is white.

2c. If Pierre has sufficient undefeated evidence that dogs bark in an epistemic circumstance *E*, then Pierre *semantically* ought to accept the sentence ‘les chiens aboient’ in *E*.

One could argue that even if Clifford could sensibly reject the inference from **1a** to **1b**, he would have to be conceptually confused to reject the inference from **1a** to **1c**. What Clifford ought to say about Ursula’s case is that she *epistemically* ought not to believe that snow is not white, but that she *semantically* ought to accept the sentence ‘schnee ist weiss’ nonetheless. Similarly, one could argue that even if Pascal could sensibly reject the inference from **2a** to **2b**, he would have to be conceptually confused to reject the inference from **2a** to **2c**. What Pascal should say is that while Pierre *prudentially* ought not to believe DOGS BARK, he *semantically* ought to accept the sentence ‘les chiens aboient’ nonetheless.

However, this response is not available to Gibbard, who claims that semantic ascriptions entail ‘Ewing’s primitive oughts’, where to say ‘you ought to do *X*’ is to say that you ought to do *X* all things considered, or that you have a conclusive reason to do *X* (Gibbard 2012:14). Thus, Gibbard’s official view is that **1a** entails **1b**, and **2a** entails **2b**, where both **1b** and **2b** are understood as all things considered oughts, not that **1a** entails **1c** nor that **2a** entails **2c**.

Moreover, dropping the suggestion that semantic oughts are Ewing oughts does not help. It is possible for someone to sensibly reject the inference from **1a** to **1c**, or from **2a** to **2c**, without displaying any conceptual confusion, so even these entailments are not plausibly analytic. To see why, it will help to generalize **1c** and **2c**. Let S be a subject; assume that sentence ‘ s ’ means P in the language spoken by S , and that ‘ s ’ is true iff p . We can generalize **1c** and **2c** as follows:

1d. S *semantically* ought to accept ‘ s ’ iff p .

2d. If S has sufficient undefeated evidence that p is true in an epistemic circumstance E , then S *semantically* ought to accept ‘ s ’ in E .

First, note that **1d** and **2d** can come into conflict.¹³ Assume that Anna speaks Swedish, in which ‘häxor existerar’ means WITCHES EXIST, it is not the case that witches exist, yet Anna has sufficient undefeated evidence that they do. In this case, **1d** will entail that it is not the case that Anna ought to accept ‘häxor existerar’ whereas **2d** will entail that S ought to accept it. Someone who holds that truth is all that matters for sentence acceptance (belief) would accept **1d** but not **2d**, whereas someone who holds that evidence is all that matters for sentence acceptance (belief), would accept **2d** but not **1d**. Yet neither of them need be confused about the concept

¹³ It is worth noting that Gibbard distinguishes between subjective and objective oughts: what you objectively ought to do is what you ought to do in light of everything that is the case, what you subjectively ought to do is constrained in some way by your state of information. He claims that semantic oughts are subjective oughts, which suggests that he would ultimately reject the implication from **1a** to **1b**. He can thereby avoid the inconsistency between the two sets of analytic entailments. But he cannot thereby avoid the objections presented above, since those turn on the thought that one might sensibly reject either entailment without displaying conceptual confusion.

MEANING. Their dispute concerns whether one ought to accept a sentence (believe something) if and only if it is true, or whether one ought to accept a sentence (believe something) only if one has sufficient undefeated evidence that it is true. It is possible for someone to endorse either position without thereby displaying confusion about the concept MEANING.

Perhaps Gibbard would argue that the dispute between the advocate of **1d** and the advocate of **2d** is a substantive normative dispute. He might suggest that the concept MEANING is normative in the sense that to grasp the concept MEANING one must accept that meaning ascriptions have *some* normative implications—one ought to accept *either 1d, 2d, or similar*—though it does not matter *which* normative implications one takes semantic ascriptions to have. But this too seems to be mistaken. For, many semantic anti-normativists will be inclined to deny that meaning ascriptions have any normative entailments. If this denial is sensible, then MEANING is not a thin normative concept. However, since I count myself among the semantic anti-normativists, I am not in a position to assess, in an impartial manner, whether our reasons for denying meaning ascriptions have any normative entailments are sensible.¹⁴

What positive reasons does Gibbard give for accepting the normativity thesis?¹⁵ The chief reason is that the normativity of the concept MEANING explains

¹⁴ See e.g. Glüer 1999, Glüer-Pagin and Wikforss 2009a, Hattiangadi 2006, 2009, 2007, Wikforss 2001. For a comprehensive review of the issue, see Glüer-Pagin and Wikforss 2009b.

¹⁵ Gibbard (2012: 16) distinguishes between a weak normativity thesis, according to which meaning ascriptions have normative entailments and a strong normativity thesis, according to which the meaning of an expression is defined by the pattern of oughts it entails. Though Gibbard purports to defend both theses, it is only the weak normativity thesis that is relevant to the present discussion.

why certain basic oughts follow from semantic ascriptions ‘invariably.’ (Gibbard 2012:16) But this rings hollow in light of the foregoing discussion. As we have seen in the cases above, there are sensible grounds to question whether semantic ascriptions really do entail basic oughts invariably.

Another reason Gibbard provides in support of the normativity thesis has to do with the hard problem of intentionality. Attempts to analyze MEANING in naturalistic terms have failed, he says, and this can be explained by the hypothesis that MEANING is normative. (Gibbard 2012: 16-17) But it would be premature to conclude that MEANING is normative on this basis. Arguably, phenomenal concepts resist reductive analysis, but it does not follow that they are normative.

The third reason Gibbard provides is that viewing the concept MEANING as normative promises a satisfactory expressivist resolution to the hard problem of intentionality. Whether this is a good reason to accept that MEANING is a normative concept depends on whether expressivism does offer a satisfactory resolution to the hard problem of intentionality. And this remains to be seen.

To assess whether expressivism offers a satisfactory resolution to the hard problem of intentionality, it will help to have a picture of how the expressivist solution to the hard problem of morality goes. The expressivist in the moral domain claims that moral statements do not describe, but express moral judgments, and that moral judgments are not straightforwardly factual beliefs. In Gibbard’s view, moral judgments are a special kind of contingency plan. Moral expressivism is attractive to those who are tempted by nihilism or naturalism. If moral statements express plans, then there is no need to postulate any normative facts out there in the world to correspond to them. This gets naturalists of the hook, for if there is no need to

postulate any normative facts, there are no normative facts crying out for naturalistic, reductive explanation. As expressivists argue, there is no need to explain what moral facts consist in. All that is needed for a fully naturalistic explanation of morality is a naturalistic explanation of ‘representation and goal-pursuit’ (Gibbard 2013:237). In other words, all that is needed for a fully naturalistic explanation of morality is a naturalistic explanation of intentionality.

This style of explanation clearly breaks down when applied to the intentional domain. First of all, the semantic expressivist’s starting point is the *semantic* thesis that the concept MEANING is normative, and hence that semantic statements do not describe, but express plans. This thesis appears to commit the expressivist at least to the existence of one semantic fact—that semantic statements do not describe, but express plans. For, if this statement itself expressed a plan, it is difficult to see what bearing the expression of this plan would have on the hard problem of intentionality. (See Hattiangadi, forthcoming)

Second, in the moral case, we have an explanation for why there is no need to give a naturalistic, reductive explanation of what goodness consists in—because we can give a semantics for moral language that does not commit us to normative facts or properties. But the parallel move in the case of intentionality does not work. The expressivist can hardly explain why there is no need for a naturalistic explanation of what meaning consists in by arguing that if the expressivist semantics for meaning ascriptions is true, then we do not need to postulate any semantic facts or properties. If semantic ascriptions express plans, as the expressivist claims, then the property of expressing a plan is a semantic property of semantic ascriptions. This property cries out for reductive naturalistic explanation.

Third, in the moral domain, expressivism is presented as hospitable to naturalism because it is assumed that it is possible to give a naturalistic explanation of intentionality. But when applied to intentionality, the expressivist turns out to be saying that we need not naturalistically explain intentionality, and that all that is needed is a naturalistic explanation of intentionality—which is blatantly inconsistent. So, expressivism does not offer a route to a satisfactory solution to the hard problem of intentionality (See Hattiangadi 2015).

Finally, it is worth noting that for the same reason that the rule following hypothesis was irrelevant to the hard problem of intentionality, the view that the concept of meaning is normative is also in a sense irrelevant to this problem.¹⁶ As we noted previously, the hard problem of intentionality has to do with the naturalistic explanation of an arbitrary representation, including normative concepts, judgments and statements. The claim that semantic concepts are normative is a hypothesis of first-order semantic theorizing about semantic concepts, on a par with semantic analyses of vague expressions, conditionals, or gradable adjectives. Of course, any reductive explanation of intentionality will have to explain whatever properties are postulated by an ideal and complete semantic theory. But that semantic theory does not directly figure in an answer to the hard problem of intentionality.

4 INTENTIONALITY AS A SOURCE OF NORMATIVITY

Let us suppose that meanings and contents are determined by constitutive rules or normative requirements of some kind. Are these rules or requirements sources of normativity? Several proponents of the normativity of intentionality have suggested

¹⁶ Thanks to John Broome for pointing this out.

that they are. For instance, Boghossian claims that what follows from the fact that you mean *green* by ‘green’ is a host of normative truths, such as that you ought to apply ‘green’ to something if and only if it is green (Boghossian 1989). Similarly, Wedgwood does not merely hold that subjects who grasp the concept of the conditional are disposed accept arguments by modus ponens; he claims further that they *ought* to do so. And Ginsborg claims not only that someone who means addition by ‘plus’ makes a certain pattern of normative judgments, she claims that these judgments are those one ought to make (see also Whiting 2009). These philosophers seem to be saying that rules (requirements or principles) that determine meaning and content are a *source* of normativity; the semantic facts engender or determine some normative facts.¹⁷ It is *because* some course of action is required by meaning or content constitutive rules that one ought to, or has a normative reason to carry out that course of action.

Following Broome (2013), we can distinguish between a strong and a weak version of the claim that some norm *R* is a source of normativity:

Strong Normativity: Necessarily, if *R* requires that you ϕ , then you ought to ϕ because *R* requires that you ϕ .

Weak Normativity: Necessarily, if *R* requires that you ϕ , then the fact that *R* requires you to ϕ gives you a normative reason (possibly *pro tanto*) to ϕ .

¹⁷ What Kathrin Glüer Pagin and Åsa Wikforss call CE normativism CITE

Are meaning or content constitutive rules either strongly or weakly normative? First consider a rule that tells you that you ought to respond with the sum of any two numbers, when asked. Of course, in many cases, responding with the sum is what you ought to do, but often when you ought to respond with the sum, it is not because responding with the sum is required by the rule that constitutes what you mean. Sometimes, you ought to respond with the sum because you will benefit—if you are sitting an examination in mathematics, for instance. At other times, you ought to respond with the sum because this is what you morally ought to do. And sometimes you ought to respond with the sum because of a desire to communicate what you believe. But it is not necessarily the case that you ought to respond with the sum whenever you are asked. Consider again the situation in which you discover that if you answer with the sum, thousands will die. Is it true that you ought to respond with the sum in this possible situation simply because this rule requires it? Surely not!

Indeed, it is implausible that meaning constitutive rules such as these are even weakly normative. The fact that what you mean by ‘plus’ requires that you respond with the sum does not necessarily give you a *pro tanto* reason to do so, as if that reason hung in the balance against the loss of life (Jarvis Thompson 2008). Of course, you often do have a reason to answer with the sum, as noted above. Such reasons could of course be present when you discover that if you answer with the sum, thousands will die. For instance, if at the moment of your discovery, you happen to be sitting a mathematics exam, you might have to weigh your own benefit against the lives of others. But in this case, it is not *because* you are semantically required to answer with the sum that you have a reason to do so; it is because doing so will benefit you. And in the absence of any such further reasons to answer with the sum,

the mere fact that answering with the sum is required by the meaning of ‘plus’ does not itself give you any reason to respond with the sum.

It might be suggested once again that there are distinct spheres of normativity: moral, prudential, epistemic, semantic, and so forth, and that distinct spheres of normativity come into conflict in the case where you are faced with a choice between answering with the sum of two numbers and allowing thousands to die, or not answering with the sum and letting them live: in this situation, what you *semantically* ought to do is respond with the sum, and what you *morally* ought to do is keep quiet. Even if what you ought to do all things considered is in this case identical with what you morally ought to do, it does not follow that it is not the case that you semantically ought to respond with the sum.

However, this point is orthogonal to the present issue. We have assumed, for the sake of argument, that meaning and content are constituted by rules or normative requirements of some kind. We have assumed for instance that if you mean addition by ‘plus’ then the semantics of your language requires that you respond with the sum of any two numbers when asked. If what you semantically ought to do is just what you ought to do in view of the relevant semantic requirements, it is trivial that you semantically ought to do what is semantically required of you. Similarly, it is trivially true that you legally ought to do what the law requires of you, that you rationally ought to do what rationality requires of you, and that you morally ought to do what morality requires of you. But even if we grant all that, there is a further question whether these systems of norms are sources of normativity, whether you ought to or have reason to do what these systems of norms require of you. To be reminded that what you semantically ought to do is what is semantically required of you does not address this further question.

Are norms of content more plausible candidates for sources of normativity? Consider Wedgwood's suggestion that if you grasp the concept of 'if' you ought to accept inferences by modus ponens. Now, it is no doubt often true that you ought to accept inferences by modus ponens. Sometimes, this is because you will benefit by doing so, and sometimes it is because others will benefit by your doing so. But even if we suppose that it is necessarily true that you ought to accept inferences by modus ponens, this is not *because* you are required to do so by the rule constitutive of the concept that you grasp.

To see why, consider someone who grasps the McGee conditional rather than the conditional—call him McGee*. Now suppose, for the sake of argument, that it is necessarily true that one ought to accept arguments by modus ponens. If this is a necessary truth, then even if McGee* grasps the McGee conditional, it will be irrational for him to reject inferences by modus ponens. The fact that McGee* ought to accept inferences by modus ponens cannot be explained by the requirement constitutive of the concept that he grasps, for that requirement tells him to reject some arguments by modus ponens. Perhaps what explains why McGee* ought to accept inferences by modus ponens is that modus ponens is a requirement of rationality. Indeed, one of the reasons Wedgwood gives for the normativity thesis is that the concept of rationality is normative for reasoning (Wedgwood CITE). But if it is rationality that is normative, then the best explanation of why one ought to accept inferences by modus ponens is that it is required by rationality, and not by conceptual requirements. Whatever the explanation why McGee* ought to accept inferences by modus ponens, it cannot be that accepting inferences by modus ponens is required by the content of the concept that he grasps. A similar objection can be made against the view that such rules are weakly normative. Even if it is supposed that necessarily,

McGee* has a reason to accept inferences by modus ponens, this cannot be explained by the requirement of the concept that he grasps.

What bearing might all this have on the hard problem of intentionality? It has been suggested that if intentionality is a source of normativity, then there will be no satisfactory solution to the hard problem of intentionality until there is a satisfactory solution to the hard problem of normativity. It has also been suggested that there can be no reductive naturalistic solution to the hard problem of normativity because normative properties are irreducible to the natural properties, and that since intentionality is a source of normativity, there can be no reductive naturalistic solution to the hard problem of intentionality either (Boghossian 1989, Wedgwood 2007). All this is true so far as it goes. But since intentionality is not a source of normativity, there are no implications for the hard problem of intentionality.

5 NORMATIVITY AND THE REDUCTION OF INTENTIONALITY

Finally, consider the view that the semantic facts are in some sense reducible to the normative facts (together with the natural facts). In what sense reducible? First, one might hold that the semantic truths can be *reductively explained* in terms of the normative and natural truths.¹⁸ Second, one might hold that the semantic facts *supervene on* the normative and natural facts. Of course, one might hold both that the semantic truths can be reductively explained in terms of the normative and the natural

¹⁸ For instance, Wedgwood claims that no account of intentionality can be given that does not mention some normative facts or properties (CITE). This suggests that normativity figures essentially in the reductive explanation of intentionality. See also Brandom CITE.

truths, and that the former supervene on the latter.¹⁹ Indeed, if normativity is essential to the reductive explanation of intentionality, this could be a reason to hold that intentionality supervenes on the normative and natural facts.

It is important to recognize that none of these normativist views is of much interest if intentionality can be reductively explained in naturalistic terms alone. For then, any normative statements added to a reductive explanation of intentionality would inessential to the adequacy of the explanation. So, we can assume for the sake of argument that intentionality cannot be reductively explained in purely naturalistic terms. This is widely regarded as one of the chief lessons to be learned from Kripke's (1982) skeptical argument, particularly his arguments to dispositionalism. Our first question then is whether appeal to some additional normative truths makes the reductive explanation of intentionality possible.²⁰

There are broadly speaking two approaches to the reductive explanation of intentionality: holistic and atomistic. The holistic approach, associated with Lewis (CITE) and Davidson (CITE), for instance, starts with the totality of relevant non-semantic information and then assigns beliefs, desires, and meanings in accordance with certain constraints on interpretation, such as most notably, some kind of

¹⁹ There is a third option: one might hold that the semantic truths can be reductively explained in terms of the normative and natural truths, but that the semantic facts supervene on the natural facts alone. We have already had occasion to consider one version of this third view: Gibbard holds that semantic concepts are normative, but that they pick out natural properties. A Cornell realist about semantic normativity might accept the same view (CITE).

²⁰ It is almost universally accepted that the normative truths supervene on the non-normative truths. Since supervenience is a transitive relation, if the semantic truths supervene on the normative truths, does it follow that they supervene on the natural truths? The answer is no, because the non-normative truths on which the normative truths supervene are not identical to the natural truths. Indeed, the non-normative truths on which the normative truths supervene might well be semantic or intentional truths, such as truths concerning the meaning of an utterance of 'I promise...' or the contents of a subject's intrinsic desires.

principle of charity. The atomistic approach, associated with Fodor (CITE), Dretske (CITE), and Millikan (CITE), for instance, starts by giving a reductive explanation of mental representations and then attempts to build a reductive explanation of beliefs, desires and the meanings of sentences on this foundation. The contents of mental representations are reductively explained in terms of causal co-variation. The simple idea is that if a mental representation X causally co-varies with the presence of Fs , then X picks out, or represents the Fs . It is well-known that this simple view faces the disjunction problem. Suppose that horses typically cause the mental representation 'horse' to be tokened,²¹ but that sometimes, such as when a horsey looking cow is viewed at a distance across a misty meadow, a cow causes 'horse' to be tokened. It follows from the simple causal co-variation thesis that 'horse' refers to either horses or horsey looking cows. To solve this problem, it is necessary to identify some property of the disposition to token 'horse' in the presence horses that is not shared with the disposition to token 'horse' in the presence of cows, and which determines that the first disposition is content-determining, while the second disposition is not (Boghossian 1989). This suggests a role for normativity in the reductive explanation of content: perhaps it is a normative property that determines whether a disposition is content-determining.

Ruth Millikan seems to endorse a view along these lines, appealing to what she refers to as a normative notion of biological function (Millikan 1984). She argues that a disposition to token a mental representation in response to a stimulus determines that the stimulus in question is the content of that representation if representing that stimulus is a biological function of the organism. For example, the

²¹ Fodor (CITE) uses 'to token' as a verb for what happens when a mental representation is instantiated in the mind in the right kind of way (i.e., in a perception or a belief).

image on a frog's retina triggered by the presence of flies represents flies because representing flies is the biological function of the image: representing flies contributed to the reproductive success of the frog's ancestors and continues to contribute to its reproductive success. Similarly, it is a biological function of the heart to pump blood because having a heart that pumps blood contributed to the reproductive success of organisms with hearts in the past, and continues to contribute to the reproductive success of organisms with hearts. If the biological function of 'horse' is to pick out the horses, then 'horse' refers to horses.

What does this have to do with normativity? Millikan suggests that biological functions are normative in the sense that they have to do with what how things are *supposed* to or *ought* to behave. The heart ought to pump blood; if it fails to do so, it malfunctions. Moreover, it is good for an organism with a heart if its heart pumps blood, bad for it when the heart ceases to do so. Similarly, if the mental representation 'horse' ought to represent horses, if it is good for the organism that tokens 'horse' that it represents horses, then horses are what 'horse' represents for that organism. And if 'horse' ought to represent horses, then this is a normative fact about that representation, determined by a normative property of the disposition to token 'horse' in the presence of horses, namely the property of being good for the organism that possesses it.

The trouble is that 'the heart ought to pump blood', despite containing the word 'ought' does not express a genuinely normative statement. What this sentence expresses is that the heart pumps blood at all possible worlds ranked highest in the context, in this case, the ordering imposed by the functioning of the organism in question (Cf. Kratzer CITE). Roughly glossed, the statement 'the heart ought to pump blood' says that the heart's pumping blood is a necessary condition for the

functioning of the organism. And this is really not a normative claim. For it is possible to coherently judge that the heart's pumping blood is a necessary condition for the functioning of the organism while judging that it ought not to be the case, in the normative sense, that the heart pumps blood. Compare: MRSA bacteria are resistant to antibiotics because they produce *penicillinase*; the biological function of producing *penicillinase* is to protect the bacteria from the harmful effects of penicillin. But the statement 'MRSA bacteria ought to produce *penicillinase*' is not a normative statement. What it is says is that producing *penicillinase* is necessary for the functioning of the MRSA bacteria. And it is perfectly coherent to judge that producing *penicillinase* is a necessary condition for the ideal functioning of MRSA bacteria while simultaneously judging in the genuinely normative sense of 'ought' that it ought not to be the case that the MRSA bacteria produce *penicillinase*. Since Millikan's theory does not offer a genuinely normative reductive explanation of intentionality, we can set it aside here.

A genuinely normative reductive explanation of intentionality (that is atomistic) would specify some genuinely normative properties of the meaning-constituting dispositions. For instance, one might take a leaf out of Wedgwood's book and say that one's *rational* dispositions determine the content of 'if'. Alternatively, one might say that the content of 'horse' is determined in part by the fact that the disposition to token 'horse' in the presence of horses is *correct*, that is, a disposition to do what one ought to do. Thus, one might explain what makes it the case that one grasps the concept of the conditional in terms of one's rational dispositions, or what makes it the case that 'horse' picks out the horses in terms of one's correct dispositions.

This type of proposal faces a dilemma. First of all, recall the distinction between what semantic, rational, prudential and all things considered normative truths. What one semantically ought to do, for instance, is what one ought to do in light of what one means, or what one ought to do to satisfy the requirements of the concept one grasps. What one rationally ought to do is what one ought to do in light of the requirements of rationality. And what one ought to do all things considered is what one ought to do in light of all relevant facts. The dilemma faced by the normativist proposal under consideration goes as follows. If semantic normative truths are included in the reductive explanation, it renders that explanation circular. If, on the other hand, only non-semantic normative truths are added to the explanation, reductive explanation fails.

First, suppose that semantic normative truths are included in the explanation, such as that one semantically ought to token 'horse' in the presence of horses, or that one semantically ought to accept inferences by modus ponens. To say that one semantically ought to token 'horse' in the presence of horses is just another way of saying that one ought to do this *in light of what one means by 'horse'*. To say that one semantically ought to accept arguments by modus ponens is to say that one ought to do this *in light of the content of the concept that one grasps*. However, if semantic terms or their equivalents occur in the reductive explanation of intentionality, that renders the explanation circular. What we wanted to know was whether the semantic truths could be explained in terms of the non-semantic truths, not whether they could be explained in terms of some semantic truths.

On the other hand, suppose that we include rational normative truths in the reductive explanation, or truths about what one ought to do all things considered. In either case, reductive explanation of intentionality fails. Suppose that accepting

arguments by modus ponens is a requirement of rationality. If rationality is a source of normativity, then it is necessarily true that one ought to accept arguments by modus ponens. This fact would not explain why someone like McGee* grasps the concept that he does. Similarly, suppose that your mental representation 'horse' picks out the horses, but that there is a possible situation in which thousands will die unless you token 'horse' in the presence of a horsey looking cow. What you ought to do all things considered is token 'horse' in the presence of the horsey looking cow, but this does not explain what makes it the case that 'horse' picks out all and only the horses. So, it seems that an atomistic reduction of intentionality is not helped by the addition of normative truths, unless they are semantically normative truths, in which case the reductive explanation that results is circular.

Let us now turn to a consideration of the holistic approach to the reduction of intentionality, such as the view of David Lewis. According to Lewis, reductive explanation of intentionality is possible only if there is an a priori entailment from the physical truths, *P* to the semantic truths, *S*. If there is an a priori entailment from *P* to *S*, then it is possible for an ideally rational being, who knows *P* to deduce *S* a priori, that is, without appeal to any further empirical information. If there is an a priori entailment from *P* to *S*, an ideal interpreter who knows all of the truths about Karl as a physical system, should be able to deduce the semantic truths about an arbitrary subject, such as Karl. The radical interpreter can accomplish this feat by appeal to a priori constraints on interpretation—such as a principle of charity—derived from our concepts of belief, desire and meaning, as defined by folk psychology (Lewis 1974).

Lewis' proposal faces several difficulties, the best known of which is the so-called 'permutation problem'. According to one version of Lewis' view, called 'global descriptivism', the radical interpreter proceeds to assign meanings to the

expressions of Karl's language in two stages, as follows. First, she (somehow) identifies the totality of sentences that Karl holds true, or would hold true under ideal conditions. This set of sentences constitutes Karl's global theory, **T**. Next, applying the principle of charity, she assigns meanings to the sentences of **T** in such a way as to render true as many sentences of that theory as is plausibly possible. The meanings of lexical expressions are meant to simply fall out of the assignment of truth-values and truth-conditions to sentences.

As many philosophers have noticed, however, it is possible to permute the meanings assigned by an interpretation to the lexical expressions of the language in such a way that they cancel each other out, so that at the level of sentences, the truth-values remain the same (Field CITE, Putnam CITE, Quine CITE, Williams CITE). The permutation problem for the global descriptivist goes as follows: from any interpretation, x , it is possible to construct a competing interpretation y that is equivalent to x in the assignment of truth values to sentences of **T**, but which assigns bizarre extensions to lexical items. For example, suppose that Karl holds true the sentence s = 'all emeralds are green.' Let x be an interpretation that assigns the set of emeralds to 'emeralds' and the set of green things to 'green', thereby rendering s true. We can easily construct an interpretation y that assigns the set of grue things to 'green' and the set of emerires to 'emeralds', where something is grue iff it is observed before t and found to be green, or is blue otherwise, and where something is an emerire iff it is observed before t and found to be an emerald, or is a sapphire otherwise. Though y is a bizarre interpretation, it too renders s true (Davidson CITE).

Lewis' solution to the permutation problem is to invoke *naturalness*. Not all properties are created equal, according to Lewis: the property of being green is more natural than the property of being grue. And this, he claims, has implications for

metasemantics. All else being equal, an interpretation of Karl which assigns the belief that all emeralds are green is better than an interpretation which assigns the belief that all emeralds are grue because the property of being green is more natural than the property of being grue. However, as Williams has recently shown, the combination of charity and naturalness do not give rise to the desired result. It is possible to construct bizarre interpretations that are at least as natural and charitable as the intended interpretation (Williams CITE. See also Hawthorne CITE).

Now, Lewis' proposal is intended to be fully reductive and fully naturalistic; he is not a normativist. But we can nevertheless ask whether the appeal to normative truths can in any way help resolve the permutation problem. Indeed, some have thought that normativity is crucial to Davidson's account of intentionality, which bears some relation to Lewis', though Davidson's aim is not to give a fully reductive explanation of intentionality (Jackman CITE, Engel CITE).²² So, we can consider the prospects of a kind of hybrid of Lewis' view and the normativist view attributed to Davidson.

Is it possible to avoid the permutation problem by adding normativity to the explanation? Perhaps we should add to Lewis' story that the radical interpreter *ought* to assign truths to the sentences of T, at least so far as is plausibly possible. Or perhaps we should say that the justification for the principle of charity lies in the normative truth that people *ought* to accept the truth. But even if these are genuinely normative truths, it does not help in any way to resolve the permutation problem. For, that problem arises after the principle of charity has been applied.

²² Like Lewis, Davidson holds that radical interpretation is possible. However, Davidson gives up on the hope of a fully reductive explanation of intentionality (Davidson 1990). Note, moreover, that it is controversial whether Davidson truly is a normativist. Cf. Glüer CITE, Davidson CITE.

Or perhaps we should say that the principles of rationality that figure in the radical interpretation of belief and desire are normative. Lewis sometimes suggests that decision theory is a refinement of folk psychology (Lewis 1974). Perhaps what Lewis should say is that it is *normative* decision theory that refines folk psychology, a theory that characterizes fundamental principles of rationality, rather than *descriptive* decision theory, which purports to describe human behavior. Then he might say that the belief that p is the state that occupies the normative belief-role; it is rationalized by evidence and in turn rationalizes behavior. So, if the radical interpreter is presumed to know which beliefs are rational in light of the subject's evidence, and which beliefs and desires rationalize the subject's behavior, she can thereby determine what the subject means.

But this would not help to resolve the permutation problem, or rather, that version of the problem that arises for assignments of belief and desire. Suppose that Karl observes a number of green emeralds, and becomes disposed to utter sentence s , that interpretation x assigns to Karl the belief that all emeralds are green, while interpretation y assigns the belief that all emeralds are grue. One might try to argue that whereas interpretation x rationalizes Karl's belief, interpretation y does not, because 'green' and 'emerald' are projectible predicates, while 'grue' and 'bleen' are not (Goodman CITE, Lewis CITE, Weatherson CITE). However, this will not do. As Davidson pointed out, it is the relation between predicates that makes them instance confirmable, as Davidson showed (1980). Though 'all emeralds are grue' may not be instance confirmable, the pair of laws 'all emeralds are grue' and 'all sapheralds are bleen' are jointly instance confirmable. Just as an observation of a green emerald confirms the hypothesis that all emeralds are green, an observation of a grue emerald confirms the hypothesis that all emeralds are grue (Davidson 1980:226).

Thus it seems that normativity does not figure essentially in the reductive explanation of intentionality. What about the view that the semantic facts supervene on the normative and natural facts? Often, the argument for supervenience turns on the possibility of reductive explanation. If intentionality cannot be reductively explained in terms of the normative and natural truths, the supervenience claim cannot be defended on this basis. Of course, it is a familiar view that reductive explanation is not needed for supervenience. For instance, so-called Type B naturalists about consciousness hold that consciousness cannot be reductively explained, but that it supervenes on the natural facts nonetheless. Similarly, one might argue that though normativity does not figure in the reductive explanation of intentionality, the semantic truths supervene on the normative and natural truths nonetheless. The trouble is that it is difficult to see what is gained by adding normative truths to the supervenience base if those normative truths do not figure essentially in the reductive explanation. A type B materialist about intentionality accepts that intentionality cannot be reductively explained in purely naturalistic terms, and holds that the intentionality supervenes on the natural nonetheless, in the service of ontological parsimony. If normative truths do not figure in the reductive explanation of intentionality, adding normative facts to the the supervenience base would appear to be an ontological extravagance.

6 CONCLUSION

In sum, we have seen normativity really has very little bearing on the hard problem of intentionality, in any of the four senses given above. The view that meaning involves rule following or a normative judgment of some kind is untenable, and in any case, has no bearing on the hard problem of intentionality. The same goes for the hypothesis that the concept of meaning is normative. In contrast, the view that

intentionality is a source of normativity would have a bearing on the hard problem of intentionality if it were true, but since it is not true, it really has no bearing on the hard problem after all. Finally, nothing is to be gained by including normative truths in the reductive explanation of intentionality, and hence we have little reason to suppose that intentionality supervenes on any normative facts or properties.

7 WORKS CITED

- Boghossian, P. 1989. 'The Rule-Following Considerations.' *Mind*, 98:507-549.
- _____. 2005. 'Rules, Meaning and Intention – Discussion.' *Philosophical Studies*, 124:185-197.
- Brandom, R. 1994. *Making it Explicit*. Cambridge MA: Harvard University Press.
- _____. 2000. *Articulating Reasons*. Cambridge MA: Harvard University Press.
- Broome, J. 2013. *Rationality Through Reasoning*. Oxford: Oxford University Press.
- Chalmers, D. 1996. *The Conscious Mind*. Oxford: Oxford University Press.
- Clifford, W. K. 1877. The ethics of belief. *Contemporary Review*. Reprinted in T. Madigan, ed. *The ethics of belief and other essays*. Amherst, MA: Prometheus, 1999: 70-96.
- Davidson, D. 1990. The structure and content of truth. *Journal of Philosophy*, 87(6):279-328.
- Gibbard, A. 1994. 'Meaning and Normativity.' *Philosophical Issues*, 5:95-115.
- _____. 1996. 'Thought, Norms, and Discursive Practice.' *Philosophy and Phenomenological Research*, 56: 699-717.

- _____. 2003. 'Thoughts and norms.' *Philosophical Issues*, 13: 83-98.
- _____. 2013. *Meaning and Normativity*. Oxford: Oxford University Press.
- Ginsborg, H. 2011. 'Primitive Normativity and Skepticism about Rules.' *The Journal of Philosophy*: 108 (5):227-254.
- _____, 2012. 'Meaning, Understanding and Normativity,' *Aristotelian Society Supplementary Volume* 86 (1): 127-146.
- Glüer, K. 1999. 'Sense and Prescriptivity,' *Acta Analytica*, 14: 111-128.
- Glüer, K. & Pagin, P. 1999. 'Rules of Meaning and Practical Reasoning.' *Synthese*, 117: 207-227.
- Glüer, I. And Wikforss, Å. 2009. 'Against Content Normativity.' *Mind*, 118: 31-70.
- _____. 2010. 'The Normativity of Meaning and Content.' *The Stanford Encyclopaedia of Philosophy* (Winter 2010 edition), Edward N. Zalta (ed.),
 URL = <http://plato.stanford.edu/archives/win2010/entries/meaning-normativity/>.
- Hattiangadi, 2003. 'Making it Implicit: Brandom on Rule Following.' *Philosophy and Phenomenological Research*, LXVI: 419-431.
- _____. 2006. 'Is Meaning Normative?' *Mind and Language*, 21: 220-240.
- _____. 2007. *Oughts and Thoughts: Rule-Following and the Normativity of Content*.
 Oxford: Oxford University Press.
- Ichikawa, J. & Jarvis, B. 2013. *The Rules of Thought*. Oxford: Oxford University Press.

- McHugh, C. 2014. Fitting belief. *Proceedings of the Aristotelian Society* 114(2):167-187.
- Millikan, Ruth (1984). *Language, Thought, and Other Biological Categories*. Cambridge, Mass.: Bradford/MIT.
- Pascal, B. 1670. *Pensées*. W.F. Trotter, trans. London: Dent, 1910.
- Peacocke, C. 1992. *A Study of Concepts*. Cambridge MA: Harvard University Press.
- Thompson, J.J. 2008. *Normativity*. Peru, IL: Open Court.
- Wedgwood, R. 2002. 'The Aim of Belief.' *Philosophical Perspectives*, 16: 276-297.
- _____. 2007. *The Nature of Normativity*. Oxford: Oxford University Press.
- _____. 2009. 'The Normativity of the Intentional.' *The Oxford Handbook of the Philosophy of Mind*. Brian P. McLaughlin and A. Bekermann, eds. Oxford: Clarendon Press, 421-36.
- Whiting, D. 2007. 'The normativity of meaning defended,' *Analysis*, 67:133-40.
- _____. 2009. Is meaning fraught with ought? *The Pacific Philosophical Quarterly*, 90(4):575-599.
- _____. 2015. Review of *Meaning and Normativity*, by Allan Gibbard. *European Journal of Philosophy*. 23(S1):e14-e18.
- Wikforss, Å. 2001. 'Semantic Normativity.' *Philosophical Studies*, 102: 203-226.
- Williamson, T. 2007. *The Philosophy of Philosophy*. Oxford: Blackwell.
- Wittgenstein, L. 1953. *Philosophical Investigations*. Oxford: Blackwell.